CLAIMS

What is claimed is:

1. A method of using a plurality of (row-identifier, value) pairs to update rows in a table of a database, the method comprising:

repeatedly finding, and storing in a structure, a block-identifier of a block that contains a row identified by a row-identifier in at least a group of (row-identifier, value) pairs, by use of a database index;

performing a vector read operation, to store in a buffer cache, a number of blocks, said blocks being identified by block-identifiers in the structure; and

repeatedly updating, in blocks in the buffer cache, each row identified in the group of (row-identifier, value) pairs, using a corresponding value in the (row-identifier, value) pairs.

2. The method of Claim 1 further comprising:

sorting the block identifiers, prior to performing the vector read operation.

3. The method of Claim 2 wherein:

the sorting is performed subsequent to storage of the block identifiers in the structure.

4. The method of Claim 1 further comprising:

subsequent to said finding and prior to said storing, checking if the block identifier has a duplicate already stored in the structure and if so then not storing the block identifier in the structure.

5. The method of Claim 1 further comprising, prior to updating:

repeating said finding of block-identifiers for all row-identifiers in the group of (row-identifier, value) pairs.

6. The method of Claim 1 wherein:

the database index is a hash index and the table is organized in a hash cluster; and

during said finding, a single directory is used to obtain the block identifier.

7. The method of Claim 1 wherein:

the database index is a B-tree index.

8. The method of Claim 1 wherein:

said structure comprises an array; and

the array has a number of entries identical to the number of blocks that can be held in the buffer cache.

9. The method of Claim 1 further comprising:

writing a plurality of logs, at least one log for each row identified in the group of (identifier, value) pairs.

10. The method of Claim 9 further comprising:

unpinning each block after updating all rows in said each block; and

flushing an unpinned block to disk only when another block needs space in the buffer cache occupied by the unpinned block.

11. The method of Claim 1 wherein:

a plurality of file offsets are provided to the vector read operation, one offset for each block in the group.

12. A carrier wave encoded with instructions to perform the method of Claim 1.

13. A computer-readable storage medium encoded with instructions to perform the method of Claim 1.

14. The computer-readable storage medium of Claim 13 being further encoded with said structure storing the block identifiers.

15. A computer comprising a processor and a memory coupled to the processor, the memory being encoded with instructions to:

automatically use a database index to look up a block identifier of a block that contains a row identified by an identifier in a plurality of (identifier, value) pairs to be used to update a table in a database;

automatically storing the block identifier in a structure in memory;

automatically repeating (using the database index to look up and storing the block identifier), for all identifiers in at least a group of (identifier, value) pairs;

automatically performing a vector read, to store in a cache, each block in a group of blocks identified by block identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

automatically modifying a row in a block stored in the cache, using a value in the plurality of (identifier, value) pairs; and

automatically repeating said modifying with each row identified in the group of (identifier, value) pairs.

16. An apparatus for using a plurality of (identifier, value) pairs to update a table of a database, each identifier in each pair identifying a row in the table, the apparatus comprising:

means for using a database index to look up a block identifier of a block that contains the row identified by an identifier in the plurality of (identifier, value) pairs;

means for storing the block identifier in a structure in memory;

means for repeating (using the database index to look up and storing the block identifier), for all identifiers in at least a group of (identifier, value) pairs;

means for performing a vector read, to store in a cache, each block in a group of blocks identified by block identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

means for modifying a row in a block stored in the cache, using a value in the plurality of (identifier, value) pairs; and

means for repeating said modifying with each row identified in the group of (identifier, value) pairs.

17. A method of using a plurality of (row-identifier, value) pairs to update a table of a database, each row-identifier in each pair identifying a row in the table, the method comprising:

finding a block-identifier of a block that contains the row identified by a row-identifier in a (row-identifier, value) pair, by use of a database index;

storing the block-identifier in a structure;

repeating (finding the block-identifier and storing the block-identifier), for all row-identifiers in at least a group of (row-identifier, value) pairs;

performing a vector read operation, to store in a buffer cache, each block in a group of blocks identified by block-identifiers stored in said structure, wherein

the group of blocks are all stored in the cache during execution of a single function call;

updating the row in the block in the cache, using the value in the (row-identifier, value) pair; and

repeating said updating with each row identified in the group of (identifier, value) pairs.